

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

PETITION FEE Under 37 CFR 1.17(f), (g) & (h) TRANSMITTAL (Fees are subject to annual revision)		Application Number	10/785,002
		Filing Date	February 25, 2004
		First Named Inventor	Akira FUJIBAYASHI
		Art Unit	
		Examiner Name	
		Attorney Docket Number	274.43165X00

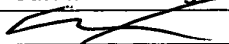
Enclosed is a petition filed under 37 CFR §1.102(d) that requires a processing fee (37 CFR 1.17(f), (g), or (h)). Payment of \$ 130.00 is enclosed.

This form should be included with the above-mentioned petition and faxed or mailed to the Office using the appropriate Mail Stop (e.g., Mail Stop Petition), if applicable. For transmittal of processing fees under 37 CFR 1.17(i), see form PTO/SB/17i.


Payment of Fees (small entity amounts are NOT available for the petition (fees))

- ☒ The Commissioner is hereby authorized to charge the following fees to Deposit Account No. 50-1417:
- ☐ petition fee under 37 CFR 1.17(f), (g) or (h) ☒ any deficiency of fees and credit of any overpayments
- Enclose a duplicative copy of this form for fee processing.
- ☐ Check in the amount of \$ _____ is enclosed.
- ☒ Payment by credit card (From PTO-2038 or equivalent enclosed). Do not provide credit card information on this form.

Petition Fees under 37 CFR 1.17(f): For petitions filed under: § 1.53(e) - to accord a filing date. § 1.57(a) - to according a filing date. § 1.182 - for decision on a question not specifically provided for. § 1.183 - to suspend the rules. § 1.378(e) for reconsideration of decision on petition refusing to accept delayed payment of maintenance fee in an expired patent. § 1.741(b) - to accord a filing date to an application under §1.740 for extension of a patent term.	Fee \$400	Fee Code 1462
Petition Fees under 37 CFR 1.17(g): For petitions filed under: §1.12 - for access to an assignment record. §1.14 - for access to an application. §1.47 - for filing by other than all the inventors or a person not the inventor. §1.59 - for expungement of information. §1.103(a) - to suspend action in an application. §1.136(b) - for review of a request for extension of time when the provisions of section 1.136(a) are not available. §1.295 - for review of refusal to publish a statutory invention registration. §1.296 - to withdraw a request for publication of a statutory invention registration filed on or after the date the notice of intent to publish issued. §1.377 - for review of decision refusing to accept and record payment of a maintenance fee filed prior to expiration of a patent. §1.550(c) - for patent owner requests for extension of time in <u>ex parte</u> reexamination proceedings. §1.956 - for patent owner requests for extension of time in <u>inter partes</u> reexamination proceedings. § 5.12 - for expedited handling of a foreign filing license. § 5.15 - for changing the scope of a license. § 5.25 - for retroactive license.	Fee \$200	Fee code 1463
Petition Fees under 37 CFR 1.17(h): For petitions filed under: §1.19(g) - to request documents in a form other than that provided in this part. §1.84 - for accepting color drawings or photographs. §1.91 - for entry of a model or exhibit. §1.102(d) - to make an application special. §1.138(c) - to expressly abandon an application to avoid publication. §1.313 - to withdraw an application from issue. §1.314 - to defer issuance of a patent.	Fee \$130	Fee Code 1464

Name (Print/Type)	Carl I. Brundidge	Registration No. (Attorney/Agent)	29,621
Signature		Date	August 4, 2005

This collection of information is required by 37 CFR 1.114. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.



274.43165X00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants: Akira FUJIBAYASHI

Serial No.: 10/785,002

Filed: February 25, 2004

For: METHOD AND APPARATUS FOR RE-SYNCHRONIZING
MIRRORING PAIR WITH DATA CONSISTENCY

PETITION TO MAKE SPECIAL
UNDER 37 CFR §1.102(MPEP §708.02)

MS Petition

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

August 4, 2005

Sir:

Applicants hereby petition the Commissioner to make the above-identified application special in accordance with 37 CFR §1.102(d). Pursuant to MPEP §708.02(VIII), Applicants state the following.

(A) This Petition is accompanied by the fee set forth in 37 CFR §1.17(h).

The Commissioner is hereby authorized to charge any additional payment due, or to credit any overpayment, to Deposit Account No. 50-1417.

(B) All claims are directed to a single invention.

If the Office determines that all claims are not directed to a single invention, Applicant will make an election without traverse as a prerequisite to the grant of special status in conformity with established telephone restriction practice.

(C) A pre-examination search has been conducted.

The search was directed towards a storage system as set forth in claims 1-28, particularly independent claims 1, 10, 15 and 24. According to claims 1, 10, 15 and 24 the invention is a method and storage system for recovering and maintaining data consistency between volumes of first and second storage systems which are connected to each other via a path, and the first storage system being connected to a host. As per the invention, a primary volume is provided to the first storage system and a remote secondary volume is provided to the second storage system, the remote secondary volume being a copied volume of the primary volume, the primary volume and the remote secondary volume being in synchronous mode, and a local secondary volume is provided in the first storage system, the local secondary volume having stored therein journal logs of write input/output (I/O) commands issued by the host to the primary volume to store write data in the primary volume and old data including write data of the write I/O commands. Recovery of data is conducted on the primary volume from the local secondary volume when necessary by performing recovery of data of the primary volume based on the old data and the journal logs while maintaining the primary volume and the remote secondary volume in the synchronous mode. Under additional aspects, according to claim 2, the journal logs are stored in a journal log table, and according to claim 3, a journal function management table is provided to be accessible and used by the first storage system to manage a process of taking journal logs as represented by the journal

log table and a process of recovery of the primary volume to a certain point in time image based on the journal logs.

The search of the above features was conducted in the following areas:

<u>Class</u>	<u>Subclasses</u>
707	201, 202, 203, 204
709	217, 219
711	100, 112, 114, 141, 161, 162, 173
714	6, 11, 20

Additionally, a computer database search was conducted on the USPTO systems EAST and WEST.

(D) The following is a list of the references deemed most closely related to the subject matter encompassed by the claims:

<u>U.S. Patent Number</u>	<u>Inventors</u>
5,555,371	Duyanovich et al
5,592,618	Micka et al
5,734,818	Kern et al
5,889,935	Ofek et al
6,308,284	LeCrone et al
6,463,501	Kern et al
6,526,419	Burton et al
6,578,120	Crockett et al

<u>U.S. Patent Application Publication No.</u>	<u>Inventor(s)</u>
2003/0135650	Kano et al

A copy of each of these references (as well as other references uncovered during the search) is enclosed in an accompanying IDS.

(E) It is submitted that the present invention is patentable over the references for the following reasons.

It is submitted that the cited references, whether taken individually or in combination with each other, fail to teach or suggest the invention as claimed. In particular, the cited references, at a minimum, fail to teach or suggest as recited in the claims:

a first feature of the present invention as recited in each of independent claims 1 and 15 of recovering data on the primary volume from the local secondary volume by performing recovery of data of the primary volume based on the old data and the journal logs while maintaining the primary volume and the remote secondary volume in the synchronous mode, and

a second feature of the present invention as recited in each of independent claims 10 and 24 of recovering data on the primary volume from the local secondary volume by performing recovery of data of the primary volume based on the base volume and the journal logs included in the local secondary volume while maintaining the primary volume and the remote secondary volume in the synchronous mode.

Further, the cited references fail to teach or suggest the above noted features of the present invention when taken in combination with the other limitations recited in the claims.

The references considered most closely related to the claimed invention are briefly discussed below:

Duyanovich (U.S. Patent No. 5,555,371) discloses a data backup copying with delayed directory updating and reduced numbers of DASD accesses at a back up site using a log structured array data storage. Primary and secondary data processing systems are coupled via a communication system. Data storage in both systems is provided by a log structured array (LSA) system that stores data in compressed form. Each time data are updated within LSA, the updated data are stored in a data storage location different from the original data. Selected data recorded in a primary data storage of the primary system is remote dual copied to the secondary system for congruent storage in a secondary data storage, such as for disaster recovery purposes. The primary systems creates a remote copying session. Within such remote copying session, a series of "pending write update" sessions are ESTABLISHED. All data updated within each pending write update session is a consistency group of data. Within each pending write update session update data are retentively stored in both the primary and secondary systems (such as in a non-volatile cache). Addressability of such update data within a pending write update session is through separate pending write directories. Congruence of the sequence of updating in the primary system is maintained in the secondary system updating procedure. The host system then, using the data mover, records the received data records in secondary data-storage system, logs the write sequence tokens in a journal and updates the pending write directory. The pending write directory can be a usual log of a host processor or be kept in the LSA. One separate log or journal can be provided for each separate pending write update session. In any event, each

copy of update data must be identifiable as being a part of a given pending write update session. (See, e.g., Abstract, column 12, lines 47-51, and column 17, lines 60-65).

However, unlike the present invention, Duyanovich does not teach or suggest a method for recovering and maintaining data consistency between volumes of first and second storage systems which are connected to each other via a path as in the present invention. Particularly, there is no teaching or suggestion in the reference that recovery of data on the primary volume is conducted from the local secondary volume when necessary by performing recovery of data of the primary volume based on the old data and the journal logs while maintaining the primary volume and the remote secondary volume in the synchronous mode as in the present invention. In addition, there is no teaching or suggestion in the reference that the recovery is conducted based on a base volume and the journal logs included in the local secondary volume while maintaining the primary and remote secondary volumes in the synchronous mode as in the present invention.

More particularly, Duyanovich at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claim 1 and 15 and the above described second feature of the present invention as recited in each of independent claims 10 and 24, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Micka (U.S. Patent No. 5,592,618) discloses a data copy validation technique that compares primary data to a copy of that data by generating check codes of the data to be validated at both primary and secondary sites in a continuously running remote data shadowing system. The validation session runs concurrently with the data shadowing so as not to disrupt the data shadowing. Each data validation session is assigned a copy start time, for example, a time preceding a time of a consistency group, for defining when the check code for data at the primary should be compared to a check code of the data at the secondary. The check code of the data at the secondary site is taken at a time that is the equivalent logical point-in-time of that data at the primary site even though the respective data contents may be different in real time. The master journal includes: consistency group number; location on journal volumes; and operational time stamp. The master journal further maintains specific record updates as grouped in consistency groups. The state table and master journal support disaster recovery, and hence must be able to operate in a stand-alone environment wherein the primary system no longer exists. (See, e.g., Abstract and column 10, line 66 to column 11, line 5).

However, unlike the present invention, Micka does not teach or suggest a method for recovering and maintaining data consistency between volumes of first and second storage systems which are connected to each other via a path as in the present invention. Particularly, there is no teaching or suggestion in the reference that recovery of data on the primary volume is conducted from the local secondary volume when necessary by performing recovery of data of the primary

volume based on the old data and the journal logs while maintaining the primary volume and the remote secondary volume in the synchronous mode as in the present invention. In addition, there is no teaching or suggestion in the reference that the recovery is conducted based on a base volume and the journal logs included in the local secondary volume while maintaining the primary and remote secondary volumes in the synchronous mode as in the present invention.

More particularly, Micka at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claim 1 and 15 and the above described second feature of the present invention as recited in each of independent claims 10 and 24, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Kern (U.S. Patent No. 5,734,818) discloses forming consistency groups using self-describing record sets for remote data duplexing. A remote data shadowing system provides storage based, real time disaster recovery capability. Record updates at a primary site cause write I/O operations in a storage subsystem therein. The write I/O operations are time stamped and the time, sequence, and physical locations of the record updates are collected in a primary data mover. The primary data mover groups sets of the record updates and associated control information based upon a predetermined time interval, the primary data mover appending a prefix header to the record (updates thereby forming self describing record sets. The self describing record sets are transmitted to a remote secondary site wherein consistency groups are formed

such that the record updates are ordered so that the record updates can be shadowed in an order consistent with the order the record updates cause write I/O operations at the primary site. The master journal includes: consistency group number; location on journal volumes; and operational time stamp. The master journal further maintains specific record updates as grouped in consistency groups. The state table and master journal support disaster recovery, and hence must be able to operate in a stand-alone environment wherein the primary system no longer exists. (See, e.g., Abstract and column 13, lines 17-23).

However, unlike the present invention, Kern '818 does not teach or suggest a method for recovering and maintaining data consistency between volumes of first and second storage systems which are connected to each other via a path as in the present invention. Particularly, there is no teaching or suggestion in the reference that recovery of data on the primary volume is conducted from the local secondary volume when necessary by performing recovery of data of the primary volume based on the old data and the journal logs while maintaining the primary volume and the remote secondary volume in the synchronous mode as in the present invention. In addition, there is no teaching or suggestion in the reference that the recovery is conducted based on a base volume and the journal logs included in the local secondary volume while maintaining the primary and remote secondary volumes in the synchronous mode as in the present invention.

More particularly, Kern '818 at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claim 1 and 15 and the above described second feature of the present invention as recited in each of independent claims 10 and 24, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Ofek (U.S. Patent No. 5,889,935) discloses disaster control features for remote data mirroring. Two data storage systems are interconnected by a data link for remote mirroring of data. Each volume of data is configured as local, primary in a remotely mirrored volume pair, or secondary in a remotely mirrored volume pair. Normally, a host computer directly accesses either a local or a primary volume, and data written to a primary volume is automatically sent over the link to a corresponding secondary volume. Each remotely mirrored volume pair can operate in a selected synchronization mode including synchronous, semi-synchronous, adaptive copy remote-write pending, and adaptive copy-disk. Each write request transmitted over the link between the data storage systems includes not only the data for at least one track in the secondary volume to be updated but also the current "invalid track" count for the secondary volume as computed by the data storage system containing the corresponding primary volume. Therefore, once a disaster occurs that destroys the data storage system containing the primary volume, the data storage system containing the secondary volume has an indication of the degree of consistency of the secondary volume. The "invalid tracks" count can be used to determine an appropriate recovery

operation for the volume, and can be used to selectively restrict read/write access to the volume when the user decides that synchronization should be required for a write access. Moreover, direct write access to a secondary volume is denied if remote mirroring is not suspended. Different application programs, for example, may have different requirements for criticality of data integrity. Certain application programs may have specific procedures, such as transaction processing or journaling facilities, for ensuring data integrity relatively independent of the data integrity of the data storage systems. (See, e.g., Abstract and column 13 and lines 18-24).

However, unlike the present invention, Ofek does not teach or suggest a method for recovering and maintaining data consistency between volumes of first and second storage systems which are connected to each other via a path as in the present invention. Particularly, there is no teaching or suggestion the reference that recovery of data on the primary volume is conducted from the local secondary volume when necessary by performing recovery of data of the primary volume based on the old data and the journal logs while maintaining the primary volume and the remote secondary volume in the synchronous mode as in the present invention. In addition, there is no teaching or suggestion in the reference that the recovery is conducted based on a base volume and the journal logs included in the local secondary volume while maintaining the primary and remote secondary volumes in the synchronous mode as in the present invention.

More particularly, Ofek at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent

claim 1 and 15 and the above described second feature of the present invention as recited in each of independent claims 10 and 24, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

LeCrone (U.S. Patent No. 6,308,284) discloses a method and apparatus for assuring data consistency in a data processing network including local and remote data storage controllers interconnected by independent communication paths. The remote storage controller or controllers normally act as a mirror for the local storage controller or controllers. If, for any reason, transfers over one of the independent communication paths is interrupted, transfers over all the independent communication paths to predefined devices in a group are suspended thereby assuring the consistency of the data at the remote storage controller or controllers. When the cause of the interruption has been corrected, the local storage controllers are able to transfer data modified since the suspension occurred to their corresponding remote storage controllers thereby to reestablish synchronism and consistency for the entire dataset. A remote storage controller mirrors a portion of the dataset in the local storage controller, such as a journal log file in a database application, over a communications path; a second communications path enables a remote storage controller to mirror the other portion of the dataset in the local storage controller, such as the database. In this configuration and with an interruption in the communications path, the remote storage controller continues to operate and mirror the corresponding dataset portion while the data in the remote storage controller remains frozen at

a point earlier in time. The database distributed over the remote storage controllers and site therefore no longer is consistent after a first write to the remote storage controller fails to be completed. (See, e.g., Abstract and column 6, lines 33-46).

However, unlike the present invention, LeCrone does not teach or suggest a method for recovering and maintaining data consistency between volumes of first and second storage systems which are connected to each other via a path as in the present invention. Particularly, there is no teaching or suggestion in the reference that recovery of data on the primary volume is conducted from the local secondary volume when necessary by performing recovery of data of the primary volume based on the old data and the journal logs while maintaining the primary volume and the remote secondary volume in the synchronous mode as in the present invention. In addition, there is no teaching or suggestion in the reference that the recovery is conducted based on a base volume and the journal logs included in the local secondary volume while maintaining the primary and remote secondary volumes in the synchronous mode as in the present invention.

More particularly, LeCrone at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claim 1 and 15 and the above described second feature of the present invention as recited in each of independent claims 10 and 24, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Kern (U.S. Patent No. 6,463,501) discloses a system, method, and program for maintaining data consistency among updates to data storage areas are provided. Each update has an update time the update was made. There are multiple groups of data storage areas. For each group, updates to the data storage area in the group are stored in a journal for storing updates to the group, wherein there are multiple journals. An indication is made in a memory area for each group of a group update time comprising a most recent update time of the updates in the group. The update time for each update in the group is not greater than the group update time. A determination is made of a minimum group update time across all the groups. At least one update is applied to storage if the update time for the update does not exceed the minimum group update time. The data storage areas may be partitioned among a plurality of computer systems, and may be maintained in at least one primary and secondary storage. The data storage areas may also comprise volumes, with each group including at least one volume. The groups of updates may also be maintained in a journal. The primary control unit initially writes data updates to a sidefile in a cache for the primary control unit. The system data mover takes the data updates from the sidefile and writes them to a journal. Within the journals, the updates are arranged into consistency groups. Consistency groups are formed for all updates to volumes that share a session. The consistency group contains records that have their order of update preserved, even across multiple storage controls. This preservation of order is absolutely vital for applications

that process dependent write Input/Output (I/Os) operations such as index and data and database and log. (See, e.g., Abstract and column 4, lines 46-57).

However, unlike the present invention, Kern '501 does not teach or suggest a method for recovering and maintaining data consistency between volumes of first and second storage systems which are connected to each other via a path as in the present invention. Particularly, there is no teaching or suggestion in the reference that recovery of data on the primary volume is conducted from the local secondary volume when necessary by performing recovery of data of the primary volume based on the old data and the journal logs while maintaining the primary volume and the remote secondary volume in the synchronous mode as in the present invention. In addition, there is no teaching or suggestion in the reference that the recovery is conducted based on a base volume and the journal logs included in the local secondary volume while maintaining the primary and remote secondary volumes in the synchronous mode as in the present invention.

More particularly, Kern '501 at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claim 1 and 15 and the above described second feature of the present invention as recited in each of independent claims 10 and 24, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Burton (U.S. Patent No. 6,526,419) discloses a method, system, program, and data structure for providing a shadow copy of data storage areas in a

primary site to data storage areas in a secondary site. A definition is made of storage areas in the primary site having data to be shadowed and corresponding storage areas in the secondary site to shadow data at the storage areas in the primary site. A shadow pair comprises one primary storage area and one secondary storage area that shadows data for the primary storage area in the pair. A standard data transfer interface protocol command, such as a SCSI or Fibre Channel command, is used to configure status storage areas in the secondary site to provide status information and data for each primary and secondary storage area. A write command in the standard data transfer interface protocol is then used to write status information to the status storage areas indicating status of the shadowed data at the secondary storage areas in the pairs. In preferred embodiments, hosts log updates in a journal until they receives status that the update was successfully applied. In the event of a failure, for those writes that have not completed, the hosts can recover the unapplied writes from the host journal log, and then apply those updates to the secondary LUNs involved in remote copy to derive the current data at the time of the failure. If multiple hosts had outstanding writes to the primary LUNs in the remote copy relationship, then the uncompleted updates would have to be gathered from journal logs from all those hosts having uncompleted updates to the primary LUNs. (See, e.g., Abstract and column 7, lines 13-24).

However, unlike the present invention, Burton does not teach or suggest a method for recovering and maintaining data consistency between volumes of first and second storage systems which are connected to each other via a path as in

the present invention. Particularly, there is no teaching or suggestion in the reference that recovery of data on the primary volume is conducted from the local secondary volume when necessary by performing recovery of data of the primary volume based on the old data and the journal logs while maintaining the primary volume and the remote secondary volume in the synchronous mode as in the present invention. In addition, there is no teaching or suggestion in the reference that the recovery is conducted based on a base volume and the journal logs included in the local secondary volume while maintaining the primary and remote secondary volumes in the synchronous mode as in the present invention.

More particularly, Burton at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claim 1 and 15 and the above described second feature of the present invention as recited in each of independent claims 10 and 24, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Crockett (U.S. Patent No. 6,578,120) discloses a method for assuring consistency between a primary volume and a remote secondary volume wherein a host CPU manages track-to-track transfers using loosely-coupled, storage control unit mediated data paths. The method includes initial volume synchronization in which concurrent updates to primary tracks are serialized by the copy progression at the secondary volume according to address. The method further includes volume resynchronization in the event of extrinsic error, fault, or the like in which bit maps and timestamps are used to determine the

status of primary tracks either lost while in flight to the secondary volume or updated during a volume suspension interval. The bit maps and timestamps preserve the most recent version copy order of the tracks on the secondary volume. More particularly, the first step of the method, namely that of initially synchronizing the tracks of data on the primary DASD volume with counterpart tracks on the secondary DASD volume, comprises several substeps. These substeps include reading from the primary volume of a predetermined number of tracks as a group in a monotonic address order and copying said tracks in that address order on the secondary volume. The substeps further include forming record sets of CPU-originated updates to the tracks on the primary volume and copying those record sets to the secondary volume having addresses less than the highest address of the primary tracks copied onto the secondary volume. (See, e.g., Abstract and column 5, lines 14-25).

However, unlike the present invention, Crockett does not teach or suggest a method for recovering and maintaining data consistency between volumes of first and second storage systems which are connected to each other via a path as in the present invention. Particularly, there is no teaching or suggestion in the reference that recovery of data on the primary volume is conducted from the local secondary volume when necessary by performing recovery of data of the primary volume based on the old data and the journal logs while maintaining the primary volume and the remote secondary volume in the synchronous mode as in the present invention. In addition, there is no teaching or suggestion in the reference that the recovery is conducted based on a base volume and the journal logs

included in the local secondary volume while maintaining the primary and remote secondary volumes in the synchronous mode as in the present invention.

More particularly, Crockett at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claim 1 and 15 and the above described second feature of the present invention as recited in each of independent claims 10 and 24, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Kano (U.S. Patent Application Publication No. 2003/0135650) discloses a backup method for a system including a network attached storage including a primary volume and a secondary volume, and a backup server connected to the network. The network attached storage performs a resynchronization process when a backup request is received at a time point that consistency of a file system can be guaranteed, the resynchronization process making the contents of the primary and secondary volumes coincide with each other, splits the secondary volume from the primary volume and transfers data of the secondary volume to the backup server while an online operation with the primary volume continues. The backup server stores the transferred data in a recording medium. The network attached storage performs again the resynchronization process for making the contents of the primary and secondary volumes coincide with each other, after backup completion. Typically, a bit map having each bit corresponding to each block of the primary volume is prepared, and the bit corresponding to the data written block is turned ON. The mirror control module

periodically copies a block or blocks corresponding to the ON bit or bits from the primary volume to the secondary volume. After this copy, the mirror control module turns OFF all bits of the bit map. An operation of making the contents of the primary and secondary volumes coincide with each other is called resynchronization. (See, e.g., Abstract and paragraph 43).

However, unlike the present invention, Kano does not teach or suggest a method for recovering and maintaining data consistency between volumes of first and second storage systems which are connected to each other via a path as in the present invention. Particularly, there is no teaching or suggestion in the reference that recovery of data on the primary volume is conducted from the local secondary volume when necessary by performing recovery of data of the primary volume based on the old data and the journal logs while maintaining the primary volume and the remote secondary volume in the synchronous mode as in the present invention. In addition, there is no teaching or suggestion in the reference that the recovery is conducted based on a base volume and the journal logs included in the local secondary volume while maintaining the primary and remote secondary volumes in the synchronous mode as in the present invention.

More particularly, Kano at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claim 1 and 15 and the above described second feature of the present invention as recited in each of independent claims 10 and 24, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Therefore, since the cited references at a minimum fail to teach or the above described first feature of the present invention as recited in each independent claims 1 and 15 and the above described second feature of the present invention as recited in each of independent claim 10 and 24, and further fail to teach or suggest these features in combination with the other limitations recited in each of the independent claims, it is submitted that all of the claims are patentable over the cited references whether said references are taken individually or in combination with each other.

(F) Conclusion

Applicant has conducted what it believes to be a reasonable search, but makes no representation that "better" or more relevant prior art does not exist. The United States Patent and Trademark Office is urged to conduct its own complete search of the prior art, and to thoroughly examine this application in view of the prior art cited herein and any other prior art that the United States Patent and Trademark Office may locate in its own independent search. Further, while Applicant has identified in good faith certain portions of each of the references listed herein in order to provide the requisite detailed discussion of how the claimed subject matter is patentable over the references, the United States Patent and Trademark Office should not limit its review to the identified portions but rather, is urged to review and consider the entirety of each reference, and not to rely solely on the identified portions when examining this application.

In view of the foregoing, Applicant requests that this Petition to Make Special be granted and that the application undergo the accelerated examination procedure set forth in MPEP 708.02 VIII.

(G) Fee (37 C.F.R. 1.17(h))

The fee required by 37 C.F.R. § 1.17(h) is to be paid by:

☒ the Credit Card Payment Form (attached) for \$130.00.

☐ charging Account _____ the sum of \$130.00.

A duplicate of this petition is attached.

Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, or credit any overpayment of fees, to the deposit account of MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C., Deposit Account No. 50-1417 (274.43165X00).

Respectfully submitted,

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.



Carl I. Brundidge
Reg. No. 29,621

CIB/jdc
(703) 684-1120